UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/675,129	09/30/2003	Matthias A. Blumrich	YOR920030236US2 (16773)	5707

7590 01/11/2008

STEVEN FISCHMAN, ESQ.
SCULLY, SCOTT, MURPHY AND PRESSER
400 Garden City Plaza
Garden City, NY 11530

EXAMINER	
CHOI, EUNSOOK	

ART UNIT	PAPER NUMBER
2619	

MAIL DATE	DELIVERY MODE
01/11/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/675,129	BLUMRICH ET AL.	
	Examiner	Art Unit	
	Eunsook Choi	2619	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *25 October 2007*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-41* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-41* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Response to Amendment

1. In the reply filed on 10/25/2007, the following has occurred:

- Claims 1, 16, 17, 18, 22, 25, 26, 32, 33, 36, 37, 38, and 39 are amended.

Claim Rejections - 35 USC § 103

2. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

3. Claims 1, 5-7, 9, 26, 29-31, and 33 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) in view of SP2 High-Performance Switch Architecture (Hot Interconnects II, 1994. Symposium Record).

Regarding claims 1 and 26, Carvey teaches in Fig. 4B And Paragraphs 16 and 51 a Gamma graph fabric requiring one hop to traverse the fabric (as in a one-bounce data network, interconnected to each other via communication links). Five of the switch ports are configured as access links and as fabric links. Carvey teaches the switch can couple to switches (a plurality of nodes) of adjacent fabric routers by links which fan into and out of the switch (the network including a plurality of interconnected switch devices). A switch can also forward at least portions of data packets to at least one traffic manager or to switches (the destination switch) of the adjacent switching nodes (a

message is communicated between any two switches passes over a single link from a source switch to a destination switch, the source switch concurrently sends a message to a switch configured for sending the message to the destination switch). It is probably inherent in Carvey, however, Carvey does not expressly teach an arbitrary bounce switch and said one-bounce network enabling simultaneous message communication between the source switch and destination switch up to the aggregate bandwidth of all the links to or from a switch. High-Performance Switch Architecture teaches in **SP2 128-way,** the 8th slide, intermediate switch boards. It would have been obvious to one of ordinary skill in the art at the time of the invention was made to have an arbitrary bounce switch and said one-bounce network enabling simultaneous message communication between the source switch and destination switch up to the aggregate bandwidth of all the links to or from a switch in order to have scalable aggregate bandwidth (slide 3 High-Performance Switch Architecture).

Regarding claims 5 and 29, Carvey and SP2 High-Performance Switch Architecture teach the limitations for claims 1 and 26 as applied above. It is inherent in High-Performance Switch Architecture a packet to be communicated includes means for indicating whether a packet has already bounced or not.

Regarding claims 6 and 30, Carvey and SP2 High-Performance Switch Architecture teach the limitations for claims 5 and 29 as applied above. Carvey teaches in Paragraph 102 and Fig. 11A The packet is received by the Traffic Manager, which appends a fabric header to the packet specifying each hop through the fabric to the packet destination as well as a hop count (a bounce bit in a packet header that is reset).

The hop count indicates the relative number of hops remaining to a binding node that allows packets to transition into the target egress based virtual network (set when the packet is bounced to said arbitrary other switch, a set bounce bit indicating that the packet can only go to the destination switch).

Regarding claims 7 and 31, Carvey and SP2 High-Performance Switch Architecture teach the limitations for claims 5 and 30 as applied above. Carvey teaches in Fig. 11A and Paragraphs 103 & 105 EgressController injects the packet into one of the managed SBT Segments corresponding to the packet source. The selected SBT Segment originates a source based tunnel to a binding node endpoint (the identified direct router is not the router of a destination node). Assuming that the fabric header specifies a path encompassing nodes 53, 54, 55, etc. (to specify a particular switch and) and a remaining hop count of 3 (specifying a direct router identifier), the packet will be injected into an SBT Segment corresponding to a remaining hop count of 3, a path corresponding to binding node 54. Carvey further teaches a packet arrives from an external link on fabric ingress port (the packet is bounced over an external link to the identified direct router)

Regarding claims 9 and 33, Carvey and SP2 High-Performance Switch Architecture teach the limitations for claims 6 and 30 as applied above. SP2 High-Performance Switch Architecture teaches inherently said bounce bit is set before injecting a packet into the network, such that said packet is guaranteed not to bounce, said packet injected for communication on an injection channel in an injection switch and delivered to the destination channel on the destination switch. Carvey teaches the

segment Identification Number provides a mechanism for EgressController to maintain the order of packets (a delivery being in order across the network).

4. Claims 22, 23, 24, and 25 are rejected under 35 U.S.C. 102(b) as being unpatentable over SP2 High-Performance Switch Architecture (Hot Interconnects II, 1994. Symposium Record) in view of Passint (US Patent 6101181).

Regarding claim 22, SP2 High-Performance Switch Architecture teaches in the 8th slide **SP2 128-way** intermediate switch boards (a one-bounce network level, one-bounce nodes with each one-bounce node having a link connecting to every other one-bounce node with each one-bounce node). However, SP2 High-Performance Switch Architecture does not expressly teach a plurality of cards, each card including 64 routers connected as a 4*4*4 torus network, with each router connected to a single node. Passint teaches in See Fig. 12 and Col. 9 Lines 50-63 a system having 128 nodes interconnected with 64 router chips in a double bristled torus topology (each router connected to a single node). There are four X dimension PC router boards (a plurality of cards) in each of four cabinets. As indicated, there are four locations within the X dimensions, four locations within the Y dimension, and four locations within the Z dimension resulting in a 4.times.4.times.4 torus topology (each card including 64 routers as a 4*4*4 torus network) It would have been obvious to one of ordinary skill in the art at the time of the invention was made to have a plurality of cards, each card including 64 routers connected as a 4*4*4 torus network, with each router connected to a single

node in order to have a high network bandwidth, a low inter-node distance, a high network bisection bandwidth and a high degree of fault tolerance (Col 1 Lines 45-49).

Regarding claim 23, SP2 High-Performance Switch Architecture and Passint teach the limitations for claim 22 as applied above. Passint teaches in Fig. 13-17, and Col. 9 Line 27-Col. 10 Line 60 four X dimension PC router boards (a plurality of cards) in each of four cabinets. There are four locations within the X dimensions, four locations within the Y dimension, and four locations within the Z dimension resulting in a 4.times.4.times.4 torus topology (a rack-level one-bounce network with each card connected by 4 rack-level links to each other card in its rack). Passint further teaches four routers are connected on the router board to form a torus connection of four routers in the X-dimension. The X-dimension does not scale beyond four connections. The four remaining ports of each router chip are connected between router chips to form the Y and Z dimensions for the torus topologies used in larger systems. Figures 13-17 show expanding the Y and Z dimensions (a rack unit comprises 16 cards).

Regarding claim 24, High-Performance Switch Architecture and Passint teach the limitations for claim 23 as applied above. Passint teaches Fig. 13-17, and Col. 9 Line 27-Col. 10 Line 60 four routers are connected on the router board to form a torus connection of four routers in the X-dimension. The X-dimension does not scale beyond four connections. The four remaining ports of each router chip are connected between router chips to form the Y and Z dimensions for the torus topologies used in larger systems. Figures 13-17 show expanding the Y and Z dimensions (comprising 64 racks,

the 64 racks in the machine being connected as a machine-level one-bounce network with each rack connected by 16 machine-level links to each other rack).

Regarding claim 25, SP2 High-Performance Switch Architecture teaches in the 8th slide **SP2 128-way** intermediate switch boards (one-bounce network). However, High-Performance Switch Architecture does not expressly teach comprising a total number of 2x routers at each level wherein x = 2(L-1), a L-level of said network having L links at each router, wherein a given level includes double the link bandwidth compared to bandwidth at a next lower level. Passint teaches in Fig. 4 and 6 two level having 4 routers, three level having 16 routers. It would have been obvious to one of ordinary skill in the art at the time of the invention was made to have a total number of 2x routers at each level wherein x = 2(L-1), a L-level of said network having L links at each router, wherein a given level includes double the link bandwidth compared to bandwidth at a next lower level in order to have a high network bandwidth, a low inter-node distance, a high network bisection bandwidth and a high degree of fault tolerance (Col 1 Lines 45-49).

5. Claims 2, 3, 8, 27, 28, and 32 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) modified by SP2 High-Performance Switch Architecture (Hot Interconnects II, 1994. Symposium Record) as applied to claims 1 and 26 above, and further in view of Passint (US Patent 6101181).

Regarding claims 2 and 27, Carvey and High-Performance Switch Architecture teach the limitations of claims 1 and 26 as applied above. Carvey teaches the switch

can couple to switches of adjacent fabric routers by links which fan into and out of the switch (all the switches in the network are paired and simultaneously communicate) in Paragraph 16 and Fig. 4B. However, Carvey and High-Performance Switch Architecture do not expressly teach each pair communicates at half the aggregate bandwidth of all the links from a switch each switch of a pair, one-half of the bandwidth serves the bounce message of the other switch of said pair. Passint teaches in in Col. 1 Lines 55-65 Bisection bandwidth is defined as the number of links that would be severed if the network were to be bisected by a plane at a place where the number of links between the two halves is a minimum. It would have been obvious to one of ordinary skill in the art at the time of the invention was made to have each pair communicates at half the aggregate bandwidth of all the links from a switch each switch of a pair, one-half of the bandwidth serves the bounce message of the other switch of said pair in order to have a high network bandwidth, a low inter-node distance, a high network bisection bandwidth and a high degree of fault tolerance (Col 1 Lines 45-49).

Regarding claims 3 and 28, Carvey, High-Performance Switch Architecture, and Passint teach the limitations of claims 2 and 27 as applied above. Passint teaches in Fig. 4 and 6 the effective bandwidth from each switch increases from one up to all the aggregate bandwidth of all the links from the switch as a number of switches comprising a group increases in number, the largest group comprising all the switches of the network.

Regarding claims 8 and 32, Carvey and High-Performance Switch Architecture teach the limitations of claims 1 and 26 as applied above. Carvey teaches increased

fabric bandwidth and scalability, while also preventing deadlock and tree saturation regardless of fabric size (provide a single deadlock-free channel across the switches rendering said network as deadlock free) (See Paragraph 68, Carvey). However, Carvey does not expressly teach independent channels for communicating message traffic internal to said switch. Passint teaches each type of virtual channel has virtual channel buffers assigned to each physical communication link and is capable of storing messages communicated between the processing element nodes over the physical communication links. The virtual channel assignment mechanism assigns an output next virtual channel number for determining the type of virtual channel to be used for routing from a next router along a given route. The next virtual channel number is assigned based on the lookup table virtual channel number and an input next virtual channel number received from a previous router along the given route (See Col 3, Lines 57-67, Passint) It would have been obvious for one having ordinary skill in the art at the time of the invention was made to use independent channels for communicating message traffic internal to said switch in order to avoid deadlock and to reduce network congestion (See Col. 3, Lines 26-28, Passint).

6. Claims 4, 10, 11, 14, 17, 18, 34, 36, and 39 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) modified by SP2 High-Performance Switch Architecture (Hot Interconnects II, 1994. Symposium Record) as applied to claims 1 and 26 above, and further in view of Jaber et al. (US Patent 7173912).

Regarding claim 4, Carvey and High-Performance Switch Architecture teach the limitations for claim 1 as applied above. However, Carvey and High-Performance Switch Architecture do not expressly teach the network is circuit-switched. Jaber, in the same field of endeavor, teaches Circuit-switch networks such as plain old telephone service (POTS) utilize transmission paths dedicated to specific users for the duration of a call and employ continuous, fixed-bandwidth transmission (See Col. 1 Lines 20-26 Jaber). It would have been obvious for one having ordinary skill in the art at the time of the invention was made to have a circuit-switched network because telecommunication networks transport voice and data according to a variety of standards and using a variety of technologies (See Col. 1 Lines 20-23 Jaber).

Regarding claims 10 and 34, Carvey and High-Performance Switch Architecture teach the limitations for claims 1 and 26 as applied above. Carvey teaches a two-dimensional torus array network which may serve as a fabric. The fabric is composed of a number of nodes, referred to as fabric routers, which are interconnected by fabric links. A fabric connects a group of external links (a switch comprises one or more routers, and external links, a router may have zero, one or more external links) (See Paragraph 3, Carvey). Jaber teaches one or more receiver-transceiver pairs (RPT) and a processing system interconnected by an internal Ethernet connection. Each RTP includes one or more internal interfaces and one or more external interfaces (See Col 5. Line 49-56, Jaber).

Regarding claim 11, Carvey, and High-Performance Switch Architecture, and Jaber teach the limitations of claims 10 as applied above. Jaber teaches optimizing

bandwidth usage (See Col. 3 Line 49-Col.4 Line 4, Jaber). Jaber further teaches an internal topology of a node within a network provides class of service (CoS) capabilities to support voice, video, and other real-time or time-sensitive applications. All IP packets are mapped to one of three priority levels as they enter the transport network. The guaranteed traffic has reserved bandwidth and is guaranteed to be transported within a defined time delay (See Col 4 Lines 5-21, Jaber).

Regarding claim 14, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claim 10 as applied above. Carvey teaches each Traffic Manager has a bus interface to the external links of the fabric exterior (serves the external links). Although there is substantial motivation for providing an external interface which conforms with some industry standard bus, the links interconnecting switches to other switches (i.e. fabric links) and the links interconnecting switches to Traffic Managers (i.e. access links) (serving the nodes of the switch) need not conform with an industry standard. The links can be optimized to incorporate mechanisms for implementing wormhole routing and avoiding tree saturation as these features are useful only on intra fabric links (See Paragraph 84, Carvey).

Regarding claim 17, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claim 10 as applied above. Carvey teaches a two dimensional torus array network which may serve as a fabric (said network is internally a multi-dimensional torus network) (See Paragraph 3, Carvey).

Regarding claim 18, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claims 10 as applied above. Carvey teaches a two dimensional

torus array network which may serve as a fabric (said network is internally a multi-dimensional switch network) (See Paragraph 3, Carvey).

Regarding claim 36, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claims 34 as applied above. Carvey teaches each Traffic Manager has a bus interface to the external links of the fabric exterior (serves the external links). Although there is substantial motivation for providing an external interface which conforms with some industry standard bus, the links interconnecting switches to other switches (i.e. fabric links) and the links interconnecting switches to Traffic Managers (i.e. access links) (serving the nodes of the switch) need not conform with an industry standard. The links can be optimized to incorporate mechanisms for implementing wormhole routing and avoiding tree saturation as these features are useful only on intra fabric links (See Paragraph 84, Carvey).

Regarding claim 39, Carvey, High-Performance Switch Architecture, Jaber teach the limitations of claims 34 as applied above. Carvey teaches a two dimensional torus array network which may serve as a fabric (said network is internally a multi-dimensional torus network) (See Paragraph 3, Carvey).

7. Claims 12, 13, and 35 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) modified by SP2 High-Performance Switch Architecture and Jaber et al. (US Patent 7173912,) as applied to claims 10 and 34 above, and in further view of O'Toole et al. (US Patent 7185077).

Regarding claim 12, Carvey, High-Performance Switch Architecture, and Jaber teach limitations for claim 10. However, Carvey, High-Performance Switch Architecture and Jaber do not expressly teach the external links have different bandwidths within a given switch. O'Toole, in the same field of endeavor, discloses the graphs generated by Georgia Tech Internet work Topology Models with bandwidth information. Links internal to the transit domains were assigned a bandwidth of 45 Mbits's, edges connecting stub networks to the transit domains were assigned 1.5 Mbits's, finally, in the local stub domain, edges were assigned 100 Mbit's (See Col. 30, Lines 43-51 O'Toole). It would have been obvious to one having ordinary skill in the art at the time of the invention was made to have different bandwidth on the external links within a switch because these are commonly used network technology: T3s, T1s, and Fast Ethernet. All measurements are averages over the five generated topologies (Fig. 1, Col. 30 Lines 43-51 O'Toole).

Regarding claim 13, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claim 10 as applied above. Jaber teaches the aggregated bandwidth for internal interfaces 104 is 4X155M and the bandwidth for each of four external interfaces 106 is 155M (matching the aggregate performance of the internal and external links) (See Fig. 4 and Col. 5 Lines 48-58 Jaber). However, Carvey and Jaber do not expressly teach the external links have different bandwidths within a given switch. O'Toole discloses the graphs generated by Georgia Tech Internet work Topology Models with bandwidth information. Links internal to the transit domains were assigned a bandwidth of 45 Mbits's, edges connecting stub networks to the transit

domains were assigned 1.5 Mbits's, finally, in the local stub domain, edges were assigned 100 Mbit's (See Col. 30, Lines 43-51 O'Toole). It would have been obvious to one having ordinary skill in the art at the time of the invention was made to have different bandwidth on the external links within a switch matching the aggregate performance of the internal and external links because these are commonly used network technology (Fig. 1, Col. 30 Lines 43-51 O'Toole) and to be able to deliver packets via high and/or low speed interfaces (See Col 6 Lines 14-55, Jaber).

Regarding claim 35, Carvey, High-Performance Switch Architecture, and Jaber teach limitations for claim 34. Jaber teaches optimizing bandwidth usage (See Col. 3 Line 49-Col.4 Line 4, Jaber). Jaber teaches an internal topology of a node within a network provides class of service (CoS) capabilities to support voice, video, and other real-time or time-sensitive applications. All IP packets are mapped to one of three priority levels as they enter the transport network. The guaranteed traffic has reserved bandwidth and is guaranteed to be transported within a defined time delay (internal links have differing bandwidths) (See Col 4 Lines 5-21, Jaber). Jaber further teaches the aggregated bandwidth for internal interfaces 104 is 4X155M and the bandwidth for each of four external interfaces 106 is 155M (matching the aggregate performance of the internal and external links) (See Fig. 4 and Col. 5 Lines 48-58 Jaber). However, Carvey and Jaber do not expressly teach the external links have different bandwidths within a given switch. O'Toole discloses the graphs generated by Georgia Tech Internet work Topology Models with bandwidth information. Links internal to the transit domains were assigned a bandwidth of 45 Mbits's, edges connecting stub networks to the transit

domains were assigned 1.5 Mbits's, finally, in the local stub domain, edges were assigned 100 Mbit's (See Col. 30, Lines 43-51 O'Toole). It would have been obvious to one having ordinary skill in the art at the time of the invention was made to have different bandwidth on the external links within a switch because these are commonly used network technology: T3s, T1s, and Fast Ethernet. All measurements are averages over the five generated topologies (Fig. 1, Col. 30 Lines 43-51 O'Toole).

8. Claims 15 and 37 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) modified by High-Performance Switch Architecture and Jaber et al. (US Patent 7173912) as applied to claims 10 and 34 above, and further in view of Shin et al. (US PGPUB 20030076833).

Regarding claims 15 and 37, Carvey and Jaber teach the limitations of claims 10 and 34. Jaber teaches the transport router 60 includes geographically distributed ports 34 connected to external routers 14. The external ports 34 form a port group 50 with point-to-multipoint connectivity between the ports 34 as externally represented by the router 80 (a first group of routers serving the external links) (See Fig. 2 and Col 5 Lines 22-29, Jaber). However, Carvey and Jaber do not expressly teach a second group of routers serving the nodes of the switch. Shin, in the same field of endeavor, teaches the steps of searching edge nodes within peer groups and defining the edge nodes reflected in internal links as ports of complex nodes, identifying parameters defining the state of each of the ports, and constructing radius, exception, and bypass for the identified parameters (group of routers serving the nodes of the switch) (See Paragraph

12 Shin). It would have been obvious to one having ordinary skill in the art at the time of the invention was made to have a first group of internal routers serving the external links and a second group of routers serving the nodes of the switch in order to select a connection path by reducing node information by use of complex node representation properly reflecting the characteristics of nodes and links within the peer groups (See Paragraph 11, Shin).

9. Claims 16 and 38 is rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) modified by High-Performance Switch Architecture and Jaber et al. (US Patent 7173912) as applied to claims 10 and 34 above, and further in view of Hardwick (US Patent 6292822).

Regarding claims 16 and 38, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claims 10 and 34. Carvey teaches each Traffic Manager has a bus interface to the external links of the fabric exterior (the external links). Although there is substantial motivation for providing an external interface which conforms with some industry standard bus, the links interconnecting switches to other switches (i.e. fabric links) and the links interconnecting switches to Traffic Managers (i.e. access links) (internal link) need not conform with an industry standard. The links can be optimized to incorporate mechanisms for implementing wormhole routing and avoiding tree saturation as these features are useful only on intra fabric links (Paragraph 84, Carvey). However, Carvey, High-Performance Switch Architecture, and Jaber do not expressly teach said network is configured to perform all-to-all messaging

each router is configured with internal links having twice the effective all-to-all bandwidth of the effective all-to-all bandwidth of the external links. Hardwick teaches for most interconnection network topologies, more bisection bandwidth is available in smaller subsections of the network than is available across the network as a whole, and latency may also be lower due to fewer hops between processors. Also, collective communication constructs in a message-passing layer typically also have a dependency on the number of processors involved. For example, barriers, reductions and scans are typically implemented as a virtual tree of processors, resulting in a latency of O(log P), while the latency of all-to-all communication constructs has a term proportional to P, corresponding to the point-to-point messages on which the construct is built (See Col 11 Lines 32-58 Hardwick). It would have been obvious to one having ordinary skill in the art at the time of the invention was made to perform all-to-all messaging, each router is configured with internal links having twice the effective all-to-all bandwidth of the effective all-to-all bandwidth of the external links in order to implement load balancing to distribute processing workload to available processors in a parallel computer (See Abstract, Hardwick).

10. Claim 19 is rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) modified by High-Performance Switch Architecture and Jaber et al. (US Patent 7173912) as applied to claim 10 above, and further in view of Passint (US Patent 6101181).

Regarding claim 19, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claim 10 as applied above. However, Carvey, High-Performance Switch Architecture, and Jaber do not expressly teaches a single level, two-level or multi-level bounce network. Passint teaches in Fig. 4 and 6 two level having 4 routers, three level having 16 routers. It would have been obvious to one of ordinary skill in the art at the time of the invention was made to have a single level, two-level or multi-level bounce network in order to have a high network bandwidth, a low inter-node distance, a high network bisection bandwidth and a high degree of fault tolerance (Col 1 Lines 45-49).

11. Claims 20, 21, 40, and 41 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carvey (US PGPUB 20020051427) modified by High-Performance Switch Architecture and Jaber et al. (US Patent 7173912) as applied to claims 10 and 34 above, and further in view of Saleh et al. (US PGPUB 20030031127).

Regarding claim 20 and 40, Carvey, High-Performance Switch Architecture, and Jaber teach the limitations of claims 10 and 34 as applied above. Carvey teaches the IngressController receiving the packet (injecting a packet on the network) selects a tunnel segment from information in the packet header and the fabric header prefixed to the packet. According to one embodiment, each EBT Segment is associated with an egress port lane identifier and a remaining hop count to the destination node serving the target egress port lane. The packet header provides the packet destination identifier, such as egress port lane EPL1, while the hop count and target destination node can be

determined from the fabric header (said router implementing means for determining a direct router on the injection switch with an external link to the destination switch based on said destination node identifier specified) (See Paragraph 90 Carvey). However, Carvey, High-Performance Switch Architecture, and Jaber do not expressly teach injecting a unicast packet. Saleh teaches a method of restoring a virtual path using dynamic unicast(See Paragraph 24 Saleh). It would have been obvious to one having ordinary skill in the art at the time of the invention was made to inject a unicast packet in order to guarantee restoration of the virtual path by creating a new physical path and provisioning the virtual path on the new physical path (See Paragraph 24 Saleh).

Regarding claims 21 and 41, Carvey, High-Performance Switch Architecture, Jaber, and Saleh teach the limitations of claims 20 and 40 as applied above. Carvey teaches the IngressController receiving the selects a tunnel segment from information in the packet header and the fabric header prefixed to the packet. According to one embodiment, each EBT Segment is associated with an egress port lane identifier and a remaining hop count to the destination node serving the target egress port lane. The packet header provides the packet destination identifier, such as egress port lane EPL1, while the hop count and target destination node can be determined from the fabric header (said direct router identifier is used within the injection switch to route a packet along internal links to said direct router having an external link to a destination switch, a packet possibly encountering other routers en route to said direct router) (See Paragraph 90 Carvey). Carvey further teaches the algorithm is guaranteed not to deadlock because buffer, FIR FIFO, and lane resources are allocated by hop number

(wherein an adaptive routing means at a router determines if a packet being routed should be bounced to another switch at each router encountered having one or more external links to other switches, and upon determining a packet is to be bounced, setting said bounce bit) (See Paragraph 188 Carvey).

Response to Arguments

12. Applicant's arguments, see pages 12-19, filed 10/25/2007, with respect to the rejection(s) of claim(s) 1, 5, 6, 7, 9, 22-26, 29, 30, and 31 under 35 U.S.C. 102 have been fully considered and are persuasive. Therefore, the rejection has been withdrawn. However, upon further consideration, a new ground(s) of rejection is made in view of High-Performance Switch Architecture.

Conclusion

13. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Eunsook Choi whose telephone number is 571-270-1822. The examiner can normally be reached on Monday-Friday 8:00-5:00 EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Chau Nguyen can be reached on 571-272-3126. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Eunsook Choi
1/7/2008

CHAU NGUYEN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2600